

Towards Research on Goal Reasoning with the TAO Sandbox

Bryan Auslander¹, Matthew Molineaux², David W. Aha¹, Allen Munro³ & Quentin Pizzini³

¹Navy Center for Applied Research in Artificial Intelligence;
Naval Research Laboratory (Code 5514); Washington, DC 20375

²Knexus Research Corporation; Springfield, VA 22153

³Rossier School of Education; University of Southern California; Redondo Beach, CA 90277
first.surname@nrl.navy.mil | matthew.molineaux@knexusresearch.com | surname@usc.edu

2009 August 11

Abstract. We describe our progress on instrumenting a Navy software simulator for use in the context of intelligent agent research. The Tactical Action Officer (TAO) Sandbox, developed at the University of Southern California, is used by officers to train for specific Navy missions. NRL and Knexus Research Corporation have integrated this simulator with intelligent agents using the Lightweight Integration and Evaluation Testbed (LIET), thus permitting the agent to play the role of a trainee. This will permit us to use the TAO Sandbox in our artificial intelligence research, where we are currently focusing on algorithms for continuous planning that can dynamically reason about what goal should be pursued at any time during a mission. This paper briefly describes our motivation for this integration, project status involving this simulator, and future goals.

1. Motivation

One of the core topics within the artificial intelligence (AI) research community is *planning*, which concerns the task of synthesizing a *plan* to accomplish a *goal* from an initial *state* using a set of *actions* (Russell & Norvig, 2003; Ghallab *et al.*, 2004). In this context, a state space S represents the entire set of possible states, an initial state $s_i \in S$ is a point in this space, a goal $G \subseteq S$ is a set of one or more points in this space, and an action $a \in A$ such that $a: S \rightarrow S$ is a mapping among states (i.e., an action permits movements among states). Thus, a planner $p(G, s_i, A)$ identifies a set of actions $a \in A$ that can potentially be executed so as to traverse from s_i to some state in $g \in G$.

Planning is an old topic dating to the earliest days of AI, and hundreds (if not thousands) of papers have been devoted to it, ranging from theoretical to applied, automated to interactive, and involving a large variety of approaches and assumptions. Also, annual international conferences are now devoted to AI planning. Practical motivations for planning abound, from industrial robotics tasks to Navy needs (e.g., controlling automated vehicles, assisting with assessing Course-of-Action plans, and providing intelligent agents – neutrals, allies, or adversaries – in training simulations). However, although a few exceptions exist (e.g., Cox, 2007), comparatively little research has been devoted to developing methods that can dynamically/automatically generate and manage their own goals, which is a desired capability of intelligent agents for continuous planning tasks (e.g., such as those involving many Navy missions).

Instrumented simulators are valuable for AI research because they can provide a source for data and evaluation metrics when integrated with intelligent agents (Molineaux & Aha, 2005). We seek a simulator that can serve as an *environment* (i.e., for accepting actions as input, executing them, and providing state information as output) for conducting continuous planning research on dynamic goal generation and management. Furthermore, we prefer to use Navy-relevant simulators in our research. In the following sections, we describe such a simulator (the TAO Sandbox) and its intended use in this context.

2. TAO Sandbox

The TAO Sandbox (TAOS) is an interactive Navy training simulator for Tactical Action Officers (TAOs) that is being developed by Allen Munro and Quentin Pizzini at USC's Center for Cognitive Technology. This system was designed with guidance from the Surface Warfare Officer's School (SWOS) in Newport,

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 11 AUG 2009		2. REPORT TYPE		3. DATES COVERED 00-00-2009 to 00-00-2009	
4. TITLE AND SUBTITLE Towards Research on Goal Reasoning with the TAO Sandbox			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Knexus Research Corp,9120 Beachway Lane,Springfield,VA,22153			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Technical Report AIC-09-155					
14. ABSTRACT We describe our progress on instrumenting a Navy software simulator for use in the context of intelligent agent research. The Tactical Action Officer (TAO) Sandbox, developed at the University of Southern California, is used by officers to train for specific Navy missions. NRL and Knexus Research Corporation have integrated this simulator with intelligent agents using the Lightweight Integration and Evaluation Testbed (LIET), thus permitting the agent to play the role of a trainee. This will permit us to use the TAO Sandbox in our artificial intelligence research where we are currently focusing on algorithms for continuous planning that can dynamically reason about what goal should be pursued at any time during a mission. This paper briefly describes our motivation for this integration, project status involving this simulator, and future goals.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 4	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

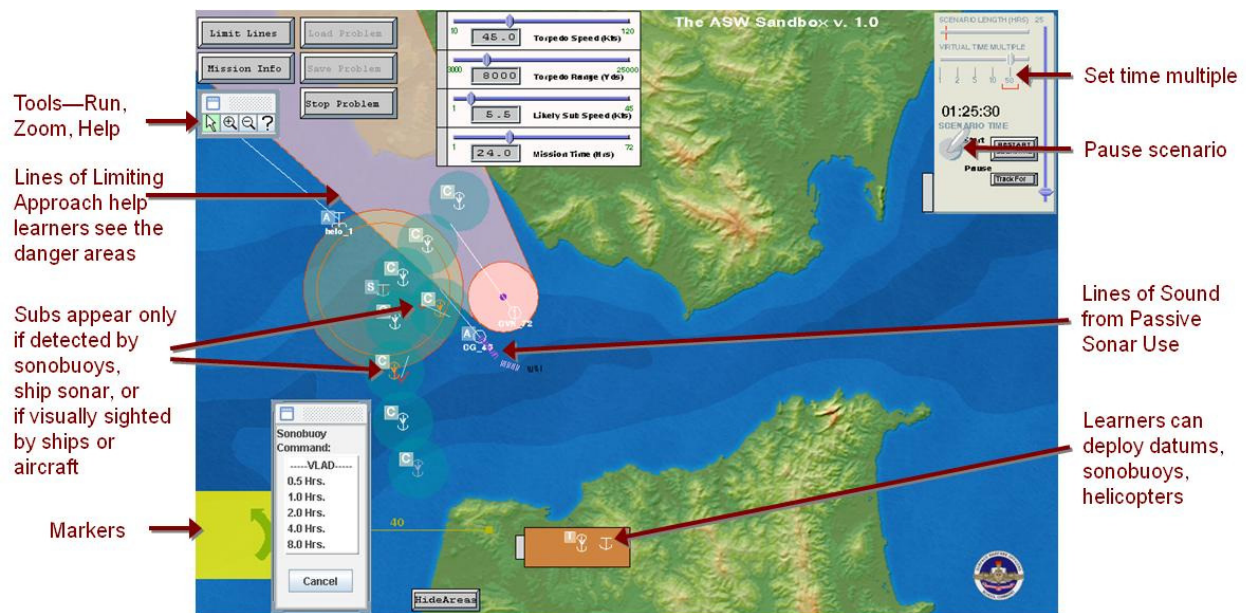


Figure 1: Snapshot of the TAO Sandbox user interface, which can be used by instructors to create simulated ASW scenarios involving the deployment and control of helicopters, sonobuoys, ships, and sonar systems. Trainee solutions provided are recorded for later review by instructors.

RI, and has been used in the SWOS Department Head Course since 2008. Trainees, working in teams, interact with the TAOS to solve conceptual problems in simulated anti-submarine scenarios (that were previously authored by instructors), during which they address decision-making situations (e.g., requiring ship coordination in a hostile environment) similar to ones that they may need to resolve later in operational situations. Trainee interactions can be recorded and replayed to assess the students' performance. The TAOS can (1) provide clear explanations of ASW issues and tactics to trainees, (2) supply more (and more varied) ASW problems to instructors, and (3) allow users to create new scenarios that may require new tactics. Figure 1 highlights aspects of the TAOS user interface and system.

A TAOS scenario includes a map containing a human-operator controlled group of ships, planes, and helicopters, along with AI-controlled enemy submarines and miscellaneous other ships that may be hostile. The operator's task is to use their units to accomplish certain mission objectives such as:

- Locate and engage enemy submarines
- Identify all ships in an area to find a target ship
- Escort a Mission Essential Unit (MEU) to a destination (e.g., a naval base)

These scenarios can be simulated to take place in any part of the world by changing or adding a map of the desired area. USC is creating more vehicle actions and adding more types of units, which further increases the system's real-world applicability and allows for more in-depth experiments.

The Navy Center for Applied Research in AI (NCARAI) at the Naval Research Laboratory is currently working with this simulator as a test-bed for current and future research. NCARAI seeks Navy simulators, such as the TAOS, for conducting AI research. TAOS is an excellent simulator to work with, due to its real-world application, use of limited knowledge (e.g., submarine locations, ship identities), and software system stability.

3. Current Research

3.1 The TAOS Interface

To date, most of NCARAI's work with the TAOS has concerned creating a software integration interface for the system. The simulator was originally designed with a human operator in mind with minimal support for automated AI. Even the computer controlled units in TAOS do not have much intelligence; they tend to travel on preset paths until they come within range of a human-controlled ship, where upon they may fire a torpedo. Given this, we created an interface for an intelligent software agent to control the operator's vehicles, and to obtain observations about the environment's state.

We designed this TAOS interface so that most necessary actions available in the system are now mapped to different methods for a controlling agent. This allows for incremental development with the TAOS and eases maintenance by being able to pinpoint computation errors. The current interface can perform all of the required actions in TAOS version 1.11. However, our interface is not entirely compatible with newer TAOS versions, which have added features (e.g., new units such as air bases, missile launchers, and surface to air missiles) and changed some functionality.

3.2 TAOS LIET Integration

We integrated the TAOS with LIET (Molineaux *et al.*, 2009), which is a free tool developed by Knexus that can be used to integrate simulation environments (such as the TAOS) with intelligent agents. LIET also provides facilities to perform experimentation and performance analysis. We used a modified version of the SHOP2 Hierarchical Task Network Planner (Nau *et al.*, 2003) and a custom-built TAO domain as the initial agent. Our modification to SHOP2 enabled *temporal* planning, which is essential for real-time environments like the one supported by TAOS simulator.

The TAOS-LIET integration has been successful. LIET can start up the TAOS, read state information from it, and convert this information into correct LISP syntax for our planning agent. LIET can also convert the plan created by this agent into actions that the TAOS simulator can execute.

3.3 Example Scenario for Dynamic Goal Reasoning and Management

TAOS offers a large playground for Dynamic Goal Reasoning and Management due primarily to the open-endedness of the goals in the system along with the ease in which one can create a custom scenario. This means that most reasonable scenarios are only limited by what objectives the researchers conceive. One simple scenario that requires dynamic goal reasoning and management is the task of identifying a ship that changes course to avoid being identified. In the simulation to identify a ship the operator needs to send a plane or a helicopter to make visual contact with an unidentified ship. In simulation terms this means the aerial unit needs to be within the visual range of the ship as defined by the current scenario. When planning for identifying a ship there is no way of knowing what the ship's course is going to be in the future. The only information available to the planner is the current speed and bearing of the ship. It is therefore possible for a ship to make an erratic course change causing the aerial vehicle not to make visual contact along its flight path. An example of this scenario is shown in Figure 2, where a helicopter was assigned to identify a ship that will soon after the assignment make a 180 degree turn, causing the helicopter to not fly into its visual range. When the plan ends or is updated the expectation of the ship being identified will not be true and dynamic reasoning or re-planning will be required to fix the situation. An explanation that the ship changed course should be generated along with a plan to fix the inconsistency. Some other simple situations that require similar dynamic reasoning are:

- While searching for an enemy submarine one of the operator controlled ships is destroyed by a torpedo.
- While identifying ships a hostile or target ship is spotted.
- A submarine sighting is reported.

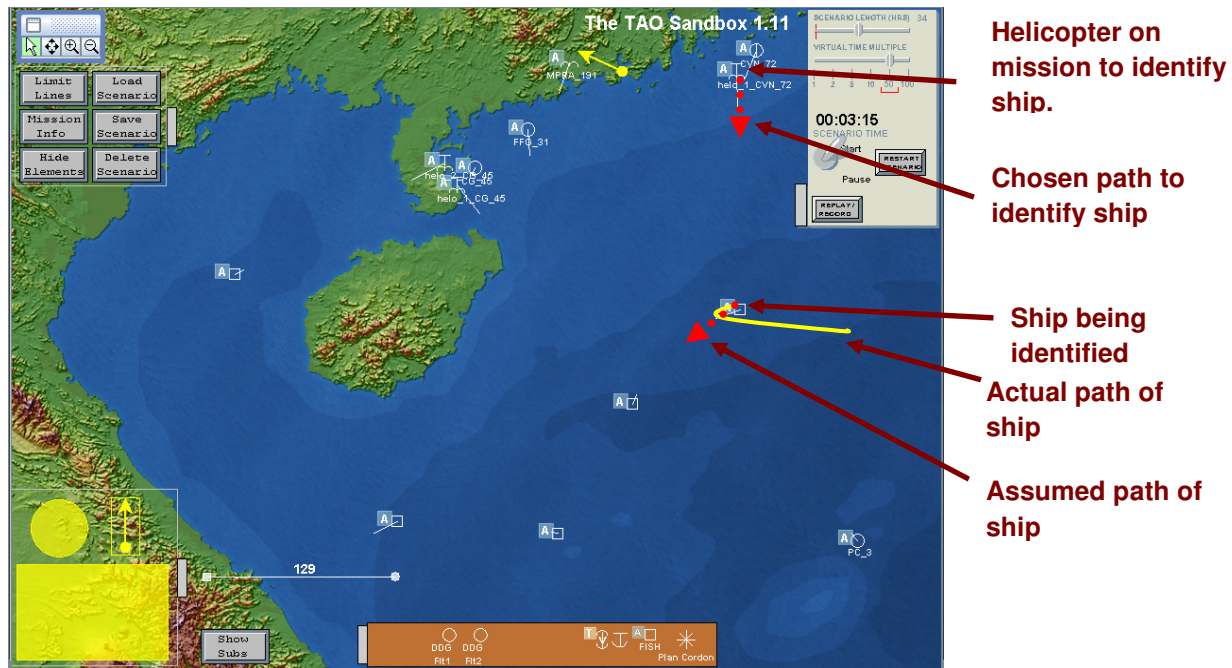


Figure 2: Snapshot of a TAO Sandbox scenario executing a plan to identify unidentified ships on the map. This scenario has as indicated by the yellow line a surprise course change for the ship being identified that will cause an aerial unit to miss it when it attempt to identify it.

4. Future Work

We integrated the TAOS with LIET to conduct research on dynamic goal reasoning and management. Through iterative system testing, we will increase the complexity of the HTN plans that our agent can process, and conduct a systematic evaluation of its components.

We also plan to use the TAOS in future research projects, including one focused on assessing the utility of relational data representations and another on studying the utility of different types of spatial-temporal representations for supporting models of computational analogy.

Bibliography

- Cox, M. T. (2007). Perpetual self-aware cognitive agents. *AI Magazine*, **28**(1), 32-45.
- Ghallab, M., Nau, D.S., & Traverso, P. (2004). *Automated planning: Theory and practice*. San Mateo, CA: Morgan Kaufmann.
- Molineaux, M., & Aha, D.W. (2005). TIELT: A testbed for gaming environments. *Proceedings of the Twentieth National Conference on Artificial Intelligence*. Pittsburgh, PA: AAAI Press.
- Molineaux, M., Aha, D.W., & Sukthankar, G. (2009). Beating the defense: Using plan recognition to inform learning agents. *Proceedings of the Twenty-Second International FLAIRS Conference* (pp. 337-343). Sanibel Island, FL: AAAI Press.
- Nau, D., Au, T.-C., Ilghami, O., Kuter, U., Murdock, J.W., Wu, D., & Yaman, F. (2003). SHOP2: An HTN planning system. *Journal of Artificial Intelligence Research*, **20**, 379-404.
- Russell, S., & Norvig, P. (2003). *Artificial intelligence: A modern approach* (2nd ed.). Upper Saddle River, NJ: Prentice Hall.